

WHAT IS CLAIMED IS:

1 1. For use in a switch fabric, a routing mechanism
2 for directing data transfers through the switch fabric
3 between an input port and an output port, wherein the
4 switch fabric includes a plurality of paths from the input
5 port to the output port, the routing mechanism comprising:
6 a controller determining, for each of the
7 plurality of paths, a likelihood of existing traffic
8 blocking a desired data transfer from the input port to the
9 output port on the respective path,

10 wherein the controller selects one of the
11 plurality of paths having a least likelihood of being
12 blocked by the existing traffic for the desired data
13 transfer.

1 2. The routing mechanism as set forth in Claim 1
2 wherein the switch fabric comprises a plurality of switches
3 each having a plurality of input ports, at least one input
4 queue associated with each input port, and a plurality of
5 output ports, wherein the input and output ports of the
6 plurality of switches are interconnected to form a multi-
7 stage switch mesh, the routing mechanism further
8 comprising:

9 a tabulation of blocking count values for each
10 output port within the multi-stage switch mesh on one of
11 the plurality of paths, wherein each blocking count value
12 represents a traffic load of existing traffic through the
13 respective output port, the controller determining the
14 likelihood of existing traffic blocking the desired data
15 transfer for each of the paths within the plurality of
16 paths by summing blocking count values for all output ports
17 on the respective path, wherein the selected one of the
18 plurality of paths has a lowest total blocking count value.

1 3. The routing mechanism as set forth in Claim 2
2 wherein the blocking count values for each output port
3 represent a traffic load of existing traffic through the
4 respective output port from both the input port for each of
5 the plurality of paths and other input ports within the
6 switch fabric.

1 4. The routing mechanism as set forth in Claim 2
2 wherein the tabulation includes a plurality of blocking
3 count values for each output port within the multi-stage
4 switch mesh on one of the plurality of paths, each of the
5 plurality of blocking count values representing a traffic
6 load of existing traffic through the respective output port
7 at a priority greater than or equal to one of a plurality
8 of priorities,

9 wherein the controller employs a blocking count
10 value corresponding to a priority of the desired data
11 transfer in determining the likelihood of existing traffic
12 blocking the desired data transfer for each of the paths
13 within the plurality of paths.

1 5. The routing mechanism as set forth in Claim 4
2 wherein each blocking count value for each output port
3 represents a traffic load of existing traffic through the
4 respective output port from any input port within the
5 switch fabric at a priority greater than or equal to a
6 corresponding one of a plurality of priorities.

1 6. The routing mechanism as set forth in Claim 2
2 wherein the tabulation includes an existing traffic load
3 for each of the plurality of paths,

4 wherein, when two or more of the plurality of
5 paths each have a same total blocking count value which is
6 lower than total blocking count values for all remaining
7 paths within the plurality of paths, the controller selects
8 one of the two or more paths having a lowest existing
9 traffic load.

1 7. The routing mechanism as set forth in Claim 6
2 wherein controller employs the existing traffic load for
3 each of the plurality of paths to identify a subset of
4 paths having sufficient capacity for the desired data
5 transfer and selects one of the subset of paths having a
6 lowest total blocking count value as a route for the
7 desired data transfer.

1 8. A switch fabric comprising:

2 a plurality of input ports;

3 a plurality of output ports, wherein each of the
4 plurality of input ports is coupled to each of the
5 plurality of output paths by two or more paths;

6 a routing mechanism for directing data transfers
7 through the switch fabric between an input port within the
8 plurality of input ports and an output port within the
9 plurality of output ports, wherein the switch fabric
10 includes a plurality of paths from the input port to the
11 output port, the routing mechanism comprising:

12 a controller determining, for each of the
13 plurality of paths, a likelihood of existing traffic
14 blocking a desired data transfer from the input port
15 to the output port on the respective path,

16 wherein the controller selects one of the
17 plurality of paths having a least likelihood of being
18 blocked by the existing traffic for the desired data
19 transfer.

1 9. The switch fabric as set forth in Claim 8 further
2 comprising:

3 a plurality of switches each having a plurality
4 of input ports, at least one input queue associated with
5 each input port, and a plurality of output ports, wherein
6 the input and output ports of the plurality of switches are
7 interconnected to form a multi-stage switch mesh, the
8 plurality of input ports for the switch fabric formed by
9 input ports for switches within a first stage of the multi-
10 stage switch mesh and the plurality of output ports for the
11 switch fabric formed by output ports for switches within a
12 last stage of the multi-stage switch mesh;

13 a tabulation of blocking count values for each
14 output port within the multi-stage switch mesh on one of
15 the plurality of paths, wherein each blocking count value
16 represents a traffic load of existing traffic through the
17 respective output port, the controller determining the
18 likelihood of existing traffic blocking the desired data
19 transfer for each of the paths within the plurality of
20 paths by summing blocking count values for all output ports
21 on the respective path, wherein the selected one of the
22 plurality of paths has a lowest total blocking count value.

1 10. The switch fabric as set forth in Claim 9 wherein
2 the blocking count values for each output port represent a
3 traffic load of existing traffic through the respective
4 output port from both the input port for each of the
5 plurality of paths and other input ports within the switch
6 fabric.

1 11. The switch fabric as set forth in Claim 9 wherein
2 the tabulation includes a plurality of blocking count
3 values for each output port within the multi-stage switch
4 mesh on one of the plurality of paths, each of the
5 plurality of blocking count values representing a traffic
6 load of existing traffic through the respective output port
7 at a priority greater than or equal to one of a plurality
8 of priorities,

9 wherein the controller employs a blocking count
10 value corresponding to a priority of the desired data
11 transfer in determining the likelihood of existing traffic
12 blocking the desired data transfer for each of the paths
13 within the plurality of paths.

1 12. The switch fabric as set forth in Claim 11
2 wherein each blocking count value for each output port
3 represents a traffic load of existing traffic through the
4 respective output port from any input port within the
5 switch fabric at a priority greater than or equal to a
6 corresponding one of a plurality of priorities.

1 13. The switch fabric as set forth in Claim 9 wherein
2 the tabulation includes an existing traffic load for each
3 of the plurality of paths,

4 wherein, when two or more of the plurality of
5 paths each have a same total blocking count value which is
6 lower than total blocking count values for all remaining
7 paths within the plurality of paths, the controller selects
8 one of the two or more paths having a lowest existing
9 traffic load.

1 14. The switch fabric as set forth in Claim 13
2 wherein controller employs the existing traffic load for
3 each of the plurality of paths to identify a subset of
4 paths having sufficient capacity for the desired data
5 transfer and selects one of the subset of paths having a
6 lowest total blocking count value as a route for the
7 desired data transfer.

1 15. For use in a switch fabric, a method of directing
2 data transfers through the switch fabric between an input
3 port and an output port, wherein the switch fabric includes
4 a plurality of paths from the input port to the output
5 port, the method comprising:

6 determining, for each of the plurality of paths,
7 a likelihood of existing traffic blocking a desired data
8 transfer from the input port to the output port on the
9 respective path; and

10 selecting one of the plurality of paths having a
11 least likelihood of being blocked by the existing traffic
12 for the desired data transfer.

1 16. The method as set forth in Claim 15 wherein the
2 switch fabric comprises a plurality of switches each having
3 a plurality of input ports, at least one input queue
4 associated with each input port, and a plurality of output
5 ports, wherein the input and output ports of the plurality
6 of switches are interconnected to form a multi-stage switch
7 mesh, the method further comprising:

8 maintaining a tabulation of blocking count values
9 for each output port within the multi-stage switch mesh on
10 one of the plurality of paths, wherein each blocking count
11 value represents a traffic load of existing traffic through
12 the respective output port; and

13 determining the likelihood of existing traffic
14 blocking the desired data transfer for each of the paths
15 within the plurality of paths by summing blocking count
16 values for all output ports on the respective path, wherein
17 the selected one of the plurality of paths has a lowest
18 total blocking count value.

1 17. The method as set forth in Claim 16 further
2 comprising:

3 setting the blocking count values for each output
4 port to represent a traffic load of existing traffic
5 through the respective output port from both the input port
6 for each of the plurality of paths and other input ports
7 within the switch fabric.

1 18. The method as set forth in Claim 16 further
2 comprising:

3 maintaining, within the tabulation, a plurality
4 of blocking count values for each output port within the
5 multi-stage switch mesh on one of the plurality of paths,
6 each of the plurality of blocking count values representing
7 a traffic load of existing traffic through the respective
8 output port at a priority greater than or equal to one of a
9 plurality of priorities; and

10 employing a blocking count value for each output
11 port corresponding to a priority of the desired data
12 transfer in determining the likelihood of existing traffic
13 blocking the desired data transfer for each of the paths
14 within the plurality of paths.

1 19. The method as set forth in Claim 18 further
2 comprising:

3 setting each blocking count value for each output
4 port to represent a traffic load of existing traffic
5 through the respective output port from any input port
6 within the switch fabric at a priority greater than or
7 equal to a corresponding one of a plurality of priorities.

1 20. The method as set forth in Claim 16 further
2 comprising:

3 maintaining, within the tabulation, an existing
4 traffic load for each of the plurality of paths;

5 employing the existing traffic load for each of
6 the plurality of paths to identify a subset of paths having
7 sufficient capacity for the desired data transfer;

8 selecting one of the subset of paths having a
9 lowest total blocking count value as a route for the
10 desired data transfer; and

11 when two or more of the subset of paths each have
12 a same total blocking count value which is lower than total
13 blocking count values for all remaining paths within the
14 subset of paths, selecting one of the two or more paths
15 having a lowest existing traffic load.